

# Doxasticism about Moral Obligation

**Kyle Kirby**

Kalamazoo College

## **ABSTRACT**

In this paper, I develop and argue for the view that an agent's moral obligations depend on their non-moral beliefs, a view which I call *doxasticism*, and I situate doxasticism within the current debate on whether an agent's moral obligations have any dependence on their epistemic or doxastic state. Three views have emerged in the contemporary literature. *Objectivism* is the view that an agent's moral obligations depend neither on their evidence nor their beliefs. *Prospectivism* is the view that an agent's moral obligations depend on their evidence. *Subjectivism* is the view that an agent's moral obligations depend on what they believe is morally best. I begin by giving a brief overview of the chief objections against each of these views. Then, I construct doxasticism from two principles: that ought implies can, and that moral obligations must be able to guide belief. In doing so, I introduce a novel modal concept of psychological possibility to describe the possibility of forming intentions to act. Lastly, I respond to two objections to doxasticism: first, that doxasticism is in conflict with robust moral realism, and second, that doxasticism unduly posits moral obligations in cases of non-veridical beliefs. I conclude that neither objection is a serious worry for the doxasticist.

## **KEYWORDS**

Moral Obligation, Belief, Objectivism, Prospectivism, Subjectivism, Ought Implies Can

## INTRODUCTION

Do an agent's moral obligations depend on their beliefs? Three central views have emerged in the contemporary literature on this question. *Objectivism* about moral obligations is the view that an agent's moral obligations depend on what is best independently of any epistemic or doxastic state of the agent. *Prospectivism* is the view that an agent's moral obligations depend on what the agent's evidence entails is best. *Subjectivism* is the view that an agent's moral obligations depend on their beliefs about what is best. In this paper, I will develop a fourth view which I call *doxasticism*, which is the view that an agent's moral obligations depend on their non-moral beliefs.

I take doxasticism to be filling a gap between the existing views. According to both subjectivism and doxasticism, an agent's moral obligations depend on some subset of their beliefs. The two views differ in what subset of their beliefs they take to be relevant; under subjectivism, the agent's moral obligations are determined by their moral opinions, while under doxasticism, the agent's moral obligations are dependent upon their non-moral beliefs.<sup>1</sup> Doxasticism is similar to both objectivism and prospectivism insofar as all of these views are non-substantive. Standing alone, none of these are complete moral theories, as each requires some additional principle about what is of moral value or about right-making in order to determine what an agent's moral obligations are. This is in contrast with subjectivism, which is a complete theory of moral obligation on its own; an agent is morally obligated to perform an action if and only if the agent believes that action to be morally best.

On one hand, objectivism and prospectivism both deny that an agent's moral obligations depend on their beliefs. Because of this feature, both of these theories face difficulties in cases of ignorance, when the agent's beliefs do not match the facts or the evidence of the case. On the other hand, subjectivism faces its own perennial problems of reducing morality to mere opinion and conflicting with robust versions of moral realism. If subjectivism is the only other theory on the market that can adequately handle cases of ignorance, then things look bleak for our theories of moral obligation. I argue that if doxasticism can capture the same

---

1. Non-moral beliefs are beliefs without moral content, which are distinct from *immoral* beliefs, or beliefs that are morally wrong to have.

insights as subjectivism in cases of ignorance while also not facing the numerous objections that subjectivism does, then we should be doxasticists.

In sections 1, 2, and 3, I briefly summarize the chief objections against objectivism, prospectivism, and subjectivism, respectively. In section 4, I construct doxasticism from two plausible and popular moral principles: that ought-implies-can, and that an agent's moral obligations must be able to provide guidance for the agent's actions. In sections 5 and 6, I respond to potential objections against doxasticism.

## **1: AGAINST OBJECTIVISM**

Objectivism is the view that an agent's moral obligations depend neither on their epistemic nor their doxastic state (Spelman 2017, 8). Another way to state this is that under objectivism, moral obligations only depend on the hard facts, where the "hard facts" are understood to be the facts unrelated to beliefs and evidence. The paradigm case against objectivism is the following example from Frank Jackson.

### *Jackson's Drug Example*

Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug. What should Jill do? (Jackson 1991, 462-463; Spelman, 9).

Any plausible objectivist account will admit of some moral principle in this case along the lines of "Jill is morally obligated to do what is best for John's well-being." Let drug C be the drug that would completely cure John. Under objectivism, Jill would be morally obligated to prescribe drug C and to not prescribe drug A. However, this seems wrong when we consider that it is generally wrong for

medical practitioners to act recklessly with regards to their patients' well-beings. By prescribing drug C, Jill is (from her point of view) making a 50/50 bet on John's life. Given that she can likely relieve John's condition with drug A without taking a massive risk on John's life, we should rather say that Jill is morally obligated to prescribe drug A instead of C.

The objectivist may respond by claiming that Jill is indeed morally obligated to prescribe drug C and not drug A, but that prescribing drug A would be an act of blameless wrongdoing. If this is the case, the objectivist needs some explanation for why intuitively, prescribing drug A not only seems blameless, but positively right. Furthermore, other cases exist in which it would not make sense to posit the objectivist obligation in the first place. Suppose that I am writing this paper while sitting in my living room, and mere feet away from me, there is a drowning child outside of my house. However, my curtains are drawn, and I cannot see the child. According to the objectivist, I would be morally obligated to walk outside my house and save the drowning child.<sup>2</sup> However, it seems much more plausible to say that I am not obligated to save the child that I have no awareness of, instead of saying that I am morally obligated to save the child but am engaging in blameless wrongdoing by continuing to write this paper. To maintain this picture of moral obligation would be to posit massive amounts of unknowable obligations and blameless wrongdoing in everyday life.

One further response the objectivist might make would be to say that these cases are different in one significant way: in Jill's case, she *can* prescribe drug C, while in my case, I *cannot* save the drowning child. Then, by appealing to ought-implies-can, the objectivist could explain why in my case, we need not posit that I had an obligation to save the child. Ultimately, this response fails because it betrays the objectivist's commitment to the principle that moral obligations only depend on the hard facts. Setting considerations about causal determinism aside, it is nomologically possible for me to save the child; no law of nature precludes me walking outside my house and picking up the child. Claiming that I could not save the child requires appealing to some sense of possibility that requires that I have awareness of the child, but the facts of my awareness are not included in the hard facts, and the objectivist cannot appeal to any such sense of possibility. These cases capture our intuitions that, as Jonathan Spelman puts it, "objectivism

---

2. I also take it that any plausible objectivist theory will generally admit of a principle such as "We are morally obligated to save drowning children instead of writing papers."

is plausible in cases where agents *know*, or at least *can know*, what is best," and implausible otherwise (Spelman, 75). This is not good enough; in searching for a general theory of obligation, we need a theory that is plausible even in cases when we aren't sufficiently informed, and so, our search continues.

## 2: AGAINST PROSPECTIVISM

Prospectivism is the view that an agent's moral obligations depend on their evidence. This time, Holly Smith provides yet another medical scenario against prospectivism.

### *Smith's Drug Example*

Harry is a physician who has to decide on the correct treatment for his patient, Renée. Careful consideration of the literature has led Harry to believe that his doing nothing (act E) is a significantly better bet than either of his alternatives, namely, prescribing drug F or drug G. Harry's senior colleague, however, knows that Harry has made a mistake. While Harry's evidence does suggest that act E is a significantly better bet than one of his alternatives (i.e., his prescribing drug F or drug G), Harry's evidence does not suggest that act E is a significantly better bet than his other alternative. In fact, Harry's evidence suggests that his other alternative is a *slightly* better bet than act E. Harry's senior colleague tells Harry this, and Harry justifiably believes her, but before he can ask her which of his alternatives is the slightly better bet, she is called away to deal with an emergency. (Smith 2010, 5; Spelman, 76).

Similarly to objectivism, I take it as a desideratum that any plausible prospectivist theory admit of a principle such as "Harry is morally obligated to do what his evidence suggests is the best bet for Renée's well-being." In this case, given that Harry's options are limited to act E, drug F, and drug G, Harry knows something about what his evidence suggests is the best bet: not-E. Either drug F or drug G is a better bet according to the evidence, but Harry does not know which one. Once again, according to prospectivism, our agent must make a coin flip. However,

being so reckless with a patient's behavior is wrong, and thus we should reject prospectivism.

Prospectivism also faces a much graver, Gettier-like problem in defining what exactly an agent's evidence is. Consider two potential answers given by Michael Zimmerman: an agent's evidence is either the evidence the agent could avail themselves of, or the evidence the agent does avail themselves of (Zimmerman 2014, 73). If an agent's evidence is the former, then there are cases like Harry's in which the agent could avail themselves of the evidence (for instance, if Harry had researched more thoroughly ahead of time) but does not, in which case the evidence can play no role in the agent's decision making. If an agent's evidence is the latter, then one could simply fail to avail themselves of any evidence for anything, thereby trivially absolving themselves of any moral obligations. In either case, how an agent's evidence is defined is too broad or too strict to be satisfactory for a general theory of obligation.

### **3: AGAINST SUBJECTIVISM**

Subjectivism is the view that an agent's moral obligations depend on what they believe is morally best. Spelman uses the cases from Jackson and Smith to motivate the case against objectivism and prospectivism. Because subjectivism is supposedly the only other theory on the market, this implicitly motivates subjectivism (Spelman 2017). The costs of moral subjectivism are generally believed to be great, but much of Spelman's paper is dedicated to showing that the costs are not as unpalatable as we may have been led to believe. He considers four arguments from Zimmerman against subjectivism and responds to them with varying degrees of success. I will be addressing the two responses I find to be least successful: his response to Zimmerman's objection that subjectivism implies that Hitler did nothing wrong, and his response to Zimmerman's objection that subjectivism implies that it is trivial to become morally infallible.

The first response is to the objection that "the Subjective View implies that, on the assumption that he was doing what he believed to be best, Hitler did no wrong. But it is grotesque to think that such a perverse belief could render mass murder morally permissible" (Zimmerman 2008. 14). Spelman's response to this case is to bite the bullet; if Hitler believed that his commanding genocide was morally best, then Hitler did not violate a moral obligation by commanding genocide. However,

Spelman claims that there are other factors that mitigate the “grotesque” nature of this claim. One consideration is that under these assumptions, while we cannot claim that Hitler did something wrong in commanding genocide, it is possible that Hitler did something wrong in forming his belief that commanding genocide was best (Spelman, 82). For this to be the case, subjectivism requires that Hitler have certain beliefs (such as the belief that we should be careful in how we form our beliefs) which in turn would give Hitler an obligation for the responsible uptake of beliefs, and we could rightly criticize Hitler for violating that obligation. Spelman does not see this as a problem; he suggests that there is already widespread agreement that we should be careful when forming our beliefs. Still, there is some small possibility that Hitler’s beliefs were arranged in such a way as never to put him in violation of a moral obligation, but Spelman believes that this would require so many layers of mitigation and unlikely scenarios that the conclusion that Hitler did nothing wrong would not be so unacceptable (Spelman, 83).

Spelman and I have opposing intuitions about whether these many layers of mitigation successfully render the conclusion acceptably not-grotesque. However, Spelman does make a point that I agree with, which is that, in extremely rare cases, when assumptions grind against many of our normal intuitions, our intuitions about the conclusions of such cases are less reliable (Spelman, 84). As it stands, I have two serious problems with Spelman’s response to this objection. The first is that one need not have intuitions about particular rare cases to hold the belief that “What Hitler did in commanding genocide was wrong, period, and it is impossible that what he did wasn’t wrong.” When considering the plausibility of that belief against belief in subjectivism, from which belief it would be better to start moral theorizing seems an open question. Philosophers who favor a top-down approach to moral theorizing, seeking to recreate our first-order intuitions like “Genocide is categorically wrong,” would see this as providing reason to reject subjectivism, rather than reason to revise our intuitions in rare cases. The second problem that I have concerns an implication of subjectivism to the Hitler case that Spelman curiously omitted. Spelman accurately notes the negative statement that in these special cases, we cannot claim that Hitler did anything morally wrong. What Spelman omits is that in the cases where Hitler believes that commanding genocide is morally best, Hitler does something positively morally right in commanding genocide. The possibility of this conclusion seems undeniably grotesque compared to the possibility of the negative statement.

The second response is to the objection that subjectivism entails that it is a trivial capability of agents to be morally infallible. It seems that we make moral mistakes all the time, so subjectivism must be mistaken. Spelman's response in defense of subjectivism is that "At times, our moral beliefs are dispositional rather than occurrent" (Spelman, 86). In these cases, given that we do not immediately know what our moral beliefs are, it does not follow that we would be able to trivially fulfill them. Furthermore, many of our intuitions about our moral mistakes stem not from the fact that we are violating our moral obligations, but instead from the fact that we are aiming to know and achieve what would be morally best, independently of our beliefs and obligations. As such, moral deliberation, solicitations of moral advice, and our intuitions that we make moral mistakes can be explained away. This defense seems to bite the bullet a little too strongly; while Spelman has provided several mitigating reasons to blunt the conclusion, it still seems quite easy to become morally infallible regarding the fulfillment of one's moral obligations.

There is also a problem with the fact that our moral beliefs are not always so clear-cut; if an agent's moral beliefs in a situation could be only dispositional, then there is a problem regarding whether those beliefs are operative in the agent's decision making. If they are, then that requires some story of how those beliefs impact the decision-making process without becoming occurrent. If they are not, then subjectivism loses an advantage that it has over objectivism and prospectivism because an agent's moral obligations are once again dependent on something that plays no direct role on their decision making. Furthermore, subjectivism might have undesirable commitments in the philosophy of belief: if one is a belief-eliminativist and a credence-realist, it is unclear how subjectivism can account for this. Credence, as a fine-grained attitude, either cannot account for the fact that we either have moral obligations or we do not, or must account for this fact with line-drawing vagueness at some credial threshold (Jackson 2020, 1). Given these outstanding objections, subjectivism also seems to be an unsatisfactory general theory of obligation.

#### **4: CONSTRUCTING DOXASTICISM**

Spelman's position is akin to David Lewis's position in *On the Plurality of Worlds*. Both saw problems with an assortment of views in the field and adopted

a theory with great costs to solve those problems. If there is a difference here, however, it is that unlike some of Lewis's opponents, we can have "paradise on the cheap:" a theory that solves our problems with objectivism and prospectivism without committing to the great costs of subjectivism (Lewis 1986, 136). Enter, doxasticism: the view that an agent's moral obligations depend on their non-moral beliefs. Like objectivism and prospectivism, doxasticism is a non-substantive moral theory: in a given case, it requires a substantive moral principle to determine exactly what an agent's moral obligations are. In Jackson's example with Jill, a plausible doxasticist principle would be "Jill is morally obligated to do what the most efficacious treatment for John's condition is according to her beliefs." Unlike subjectivism, which would require Jill to have a moral opinion in order for her to have a moral obligation to prescribe drug A, doxasticism can accommodate our intuitions in this example with only Jill's belief that drug A is the most efficacious drug.

The two Drug Examples that I used as motivation against objectivism and prospectivism share a common quality: each hinge upon the agent in question lacking crucial information, which leads us to the conclusion that their moral obligation cannot stem from that information, whether it be information about the hard facts of the case or information about what their evidence entails. I believe that the intuitions which led us to those conclusions can be explained by deeper principles of normative ethics. As such, my aim is not to further develop a top-down account of doxasticism by presenting more cases and finding a theory that fits our intuitions. Rather, my goal is to develop a bottom-up account of doxasticism by constructing the mild claim that moral obligations depend on non-moral beliefs from two plausible and popular principles in normative ethics: that ought-implies-can, and that moral obligations must be able to guide belief. Some consequences follow from this method. It may be that the principles I use are false and thus my construction fails, or that I could make a stronger claim by appealing to more or different principles. I do not believe it necessary to further address these possibilities here.

Ought-implies-can is not one single principle, but rather a family of principles, not all of which even relate to morality (King 2019, 8).<sup>3</sup> While there are many senses of "ought," the relevant one for the doxasticist is the sense of "ought" that pertains to an agent having a moral obligation. As for the other term in the

---

3. See Motiz (2012) and Wedgewood (2013) for non-moral conceptions of ought-implies-can.

implication, “can” is generally understood to be a modal notion; in some sense, a subject S can do action A if it is possible for S to do A. Typically, this form of possibility has been characterized as nomological possibility; in Kant’s formulation, the principle is stated as “The action to which the “ought” applies must indeed be possible under natural conditions,” where possibility under natural conditions is definitionally equivalent to nomological possibility, or possibility given the laws of nature (Kant 2007, A548/B576). Nomological possibility is not a quality restricted to actions; while the objects of moral obligations are actions, in statements like “It is possible that it will rain tomorrow,” nomological possibility can be applied to non-action states of affairs or propositions.

Nomological possibility, as far as kinds of possibility go, is quite narrow, especially when contrasted with logical or metaphysical possibility. However, when we discuss actions, it seems that there is an even narrower form of possibility to which we can appeal. Recall the previous example of the drowning child outside of my house that I am unaware of. Earlier, I alluded to some sense of possibility according to which it is not possible for me to save the drowning child. While that sense of possibility is not available for the objectivist, it is available for the doxasticist. This form of possibility is what I will call psychological possibility.

If nomological possibility is the possibility of a state of affairs given the laws of nature, we can characterize psychological possibility as the possibility of actions given the laws of *human* nature; or in this case, the laws of psychology. This is a mere characterization though; I do not intend to claim that there is such a thing as human nature. Rather, I wish to appeal to certain facts concerning the psychology and philosophy of action that are relevant to our discussion of moral obligations. First, we must understand the things we are morally obligated to do. We are morally obligated to do actions, but what exactly are those? One direction in which we can search for answers lies in the philosophy of language within speech act theory. In rebuffing a hypothetical person skeptical of the claim that linguistic communication involves acts, John Searle makes the following case:

...when he takes a noise or a mark on paper to be an instance of linguistic communication, as a message, one of the things that is involved in his so taking that noise or mark is that he should regard it as having been produced by a being with certain intentions. He cannot just regard it as a natural phenomenon, like a stone, a waterfall, or a tree. In order to regard it as an instance of linguistic

communication one must suppose that its production is what I am calling a speech act. (Searle 2013, 222)

In this passage, Searle makes a critical observation: what makes an action an action is not simply a physical property, but the logical presupposition of some intention behind the action. There is, admittedly, a wider sense of actions in which unintentional acts could be called actions. We could say that sleepwalking is an action, or that our involuntary heart beating is an action of the heart, or that we can anthropomorphize non-living things as taking actions, like a volcano erupting. However, this wider sense of action is not relevant for a discussion of moral obligation; these unintentional or non-living actions are simply not the kinds of things to which we ascribe moral qualities. As far as we are presently concerned, the actions we are interested in are intentional actions.

Once we grant that the objects of moral obligations are intentional actions, it follows that these actions require particular beliefs. In the earlier drowning child case, it was not psychologically possible for me to save the child because I did not have the requisite kind of beliefs to form the intention to act to save the child; namely, I was completely unaware of the child's existence at all. Therefore, I had no moral obligation to save the child- because my having an obligation to save the child would have entailed that it was psychologically possible to save the child, which would have required that I have the requisite beliefs such that I could have formed the intention to save the child. As such, there is an *awareness* condition on an agent's beliefs for their having a moral obligation: the agent must have certain descriptive beliefs such that it is possible for them to form the intention to perform the action of the obligation. We may formalize psychological possibility as the following:

*Psychological Possibility*

An action is psychologically possible if and only if it is nomologically possible for the agent to form the intention to perform the action.

Given this definition, the principle of "ought implies can" is the condition that for an agent to be morally obligated to perform an action, it must be psychologically possible for the agent to perform that action.

This condition alone is insufficient for explaining our intuitions in the Drug Examples. Given what I have stipulated, it is still psychologically possible for Jill to prescribe drug C and for Harry to prescribe drug G; it might appear bizarre or irrational to an observer, but there is no lack of awareness of the possibility of prescribing either of those drugs. Rather, what Jill and Harry lack are the beliefs about which drugs would be most effective (or in Harry's case, what the evidence entails about each of the drugs). The next principle to which I will appeal to explain our intuitions in the Drug Examples is that moral obligations must be able to guide action. How I will understand this principle is that for moral obligations to be able to guide action, when an agent has a moral obligation, some aspect of the obligation must provide a basis that is available to the agent and that upon reflection provides reasons to act. What could this basis be? With a modification to the first Drug Example, we can rule out the fact of the obligation itself serving as that basis.

### *Forgetful Jill*

While Jill is deliberating over which drug to prescribe John, she suddenly remembers a past case like John's, and recalls that she justifiably concluded that the correct drug was drug C. However, upon this realization, Jill consults the available evidence once again, and still cannot find any evidence that would inform her of which drug between B and C cured or killed; she seemingly cannot find the evidence she remembers having previously, nor does she remember what the evidence exactly is. What should Jill do?

In this example, past-Jill is a stand-in for the objectivist case, as Jill is still unaware of which drug is the completely curing drug. Even if Jill is right, and she accurately remembers prescribing drug C previously, to prescribe drug C in this case when she would have no evidential basis for prescribing C over B would still be a great risk to John's life. Without knowing why she would be obligated to prescribe drug C, above and beyond knowing the fact that she could be obligated to prescribe drug C, Jill's presumed obligation to prescribe drug C fails to provide her with adequate guidance, and therefore must not be her actual obligation at all. Furthermore, if the fact of the obligation must be available as a reflective basis

whenever an agent has a moral obligation, then it seemingly becomes impossible not to be aware of one's moral obligations. However, we frequently do not know what our moral obligations are, so this must be incorrect.

My proposal here is that moral obligations must be able to guide action in the sense that the agent has some beliefs in alignment with the explanation of the moral obligation; in other words, the agent must have some ability to know why they have that obligation. Daniel Fogal and Olle Risberg provide a recipe for explanations of particular moral facts:

*Particular descriptive explanans:* particular natural fact(s) (e.g. *a* is a lie).

*General moral explanans:* general explanatory moral principle (e.g. lying is wrong).

---

***Particular moral fact:* particular moral fact (e.g. *a* is wrong). (Fogal and Risberg 2020, 175)**

This tripartite model of moral explanations accounts for the supervenience of moral properties on natural properties and applies to talk of moral obligations, as the property of having a moral obligation is a moral property. It works through the application of general moral principles to particular natural facts, and the two function together to explain a particular moral fact. The primary goal of moral explanations is to explain *why* particular moral facts obtain, as opposed to merely stating that they do. If there is no necessary link between an agent having a moral obligation and an agent having some awareness of why they would have such an obligation, then all fulfillment of moral obligation would essentially be a matter of luck. In the above case, the explanation for why an agent would be morally obligated to save the child follows the same recipe:

*Particular descriptive explanans:* a child is drowning.

*General moral explanans:* agents have moral obligations not to let children drown.

***Particular moral fact: I have a moral obligation not to let the child drown.***

However, the fact that I do not know that the child is drowning is a matter of luck. If I am indeed still obligated to save the drowning child, there is no quality or quantity of moral reasoning that will allow me to fulfill my obligation, as I lack the beliefs that would allow the relevant moral principle (that we ought not let children drown) to fit my descriptive understanding of the situation. Thus, we arrive at the undesirable conclusion: if access to the moral explanation is not necessary for an agent to have a moral obligation, then it seems that there is no function for moral reasoning or justification doing the right thing.

So far, I have shown that at least some awareness of the moral explanation for a particular moral obligation is necessary for an agent to have that obligation. In each of our examples so far, the awareness that the agent lacks is the descriptive awareness, rather than the general moral principle. Jill presumably possesses the awareness that what would make her action right is acting for John's well-being, and I certainly know that saving drowning children is right. The reason our examples are all instances of descriptive ignorance rather than moral ignorance is because the requisite awareness of the moral explanation required for having a moral obligation is the descriptive rather than the moral explanans. To illustrate, consider the following case in which the agent has the appropriate descriptive beliefs but lacks the moral beliefs.

*Egoist Earl*

Earl lives in a world where the correct moral right-maker is the principle of utility. After reading "Famine, Affluence, and Morality" by Peter Singer in Earl's moral philosophy class, Earl learns of the descriptive facts regarding human suffering around the world. However, Earl's current moral beliefs dispose him towards an egoist view of right-making, and as such, Earl does not donate to any effective altruism causes. (Singer 1972)

What are we to make of Earl's moral obligations? Given that we have stipulated that subjectivism is not the correct right-maker in Earl's world, it seems that any observer in Earl's world would be warranted in blaming Earl for violating his moral

obligation. Earl's moral obligation only required him to have certain descriptive beliefs for the general right-making principle of utility to apply. The case could be changed such that Earl never did learn of the descriptive facts surrounding suffering in the world, but note that in that case, we would not say that Earl has a moral obligation to donate to an effective altruism cause, just as we did not say that I had a moral obligation to save the drowning child of which I was unaware.

Given the scope of our arguments from psychological possibility, we may claim that descriptive beliefs can affect what moral obligations we have, but we cannot make the parallel claim for moral beliefs. This results from the fact that descriptive beliefs carry modal force that bears on psychological possibility, while moral beliefs do not. Remember that psychological possibility requires the nomological possibility of the formation of an intention for an action. What intentions it is nomologically possible for us to form is a matter of our descriptive beliefs rather than our moral beliefs, and as such, it is the descriptive component of the moral explanation that an agent must be aware of rather than the moral component. On the one hand, if I believe that there is a wall in front of me, and I have auxiliary beliefs about the physical world such that I do not believe that I can walk through physical objects, then it seems nomologically impossible for my beliefs cause the intention to walk through the wall. On the other hand, if I believe that a course of action is merely wrong rather than physically impossible, it is much more conceivable that it is nomologically possible to intend that action anyways, regardless of the fact that I likely won't. Moral beliefs simply lack the modal force on psychological possibility that descriptive beliefs possess; perhaps if there was an instance of some anomalous person who was causally determined to always form intentions in accordance with what they believed to be morally right, this would need revision, but no such counterexample seems to exist. Furthermore, such an example would make moral beliefs relevant qua their arbitrary specification in some anomalous causal principle, rather than qua being moral beliefs.

In addition to our earlier awareness condition on an agent's beliefs for their having a moral obligation, we now have an additional guidance condition. For an agent to have a particular moral obligation, the agent must have beliefs that would fit the descriptive explanans that partially explains why they have that obligation. This is the dependence that I claim doxasticism, the view that moral obligations depend on our non-moral beliefs, captures. What substantive principle is true, if any, is a much larger question for which the bulk of normative ethics is responsible.

However, the agent need not possess the beliefs about the moral principle in order to have a particular moral obligation that arises from it. Instead, the agent only must possess the descriptive beliefs so that the principle can apply and explain the agent's moral obligation. This view is highly advantageous as it captures the informational dependence of obligations that subjectivism accounts for, while not making the further claim that morality is a matter of an agent's moral opinion. Restricting the dependence of moral obligations to non-moral beliefs side-steps the kind of relativization that makes subjectivism unappealing while preserving our original intuitions that led us to reject objectivism and prospectivism. In this way, doxasticism is paradise on the cheap.

## 5: OBJECTION FROM ROBUST REALISM

Classical theories of moral subjectivism are often excluded from the label of "moral realism," despite the fact that under subjectivism, there are still moral statements that are truth-apt and indeed, true. The belief-dependence of moral facts under many subjectivist theories may allow them to be considered "realist," but only in a weak sense. This weak sense of realism is contrasted with what David Enoch calls "robust realism," which is his label for metaethical non-naturalist views that see moral truth as "perfectly objective, universal, absolute" (Enoch 2011, 1). Given the great amount of progress recently made in advancing metaethical non-naturalism, it would be a serious concern if the belief-dependence of moral facts under doxasticism was incompatible with robust moral realism.

Luckily, Enoch provides a potential solution for the compatibility of doxasticism and robust realism. For many moral realists, only a certain kind of mind-dependence is problematic (Hanson 2018, 47). Enoch provides a criterion of normative relevance for determining which kinds of mind-dependence are acceptable: a normative truth is acceptably mind-dependent if it is only mind-dependent to the extent that the mind-dependence is normatively relevant (Enoch, 3-4). Furthermore, the normative truths must themselves not be constitutively mind-dependent, but rather must be mind-dependent according to deeper normative truths that themselves are not mind-dependent. Doxasticism fits this criterion: it specifies non-moral beliefs as normatively relevant, and the truth of doxasticism itself would not be mind-dependent if robust moral realism is true.<sup>4</sup>

---

4. There may be views of "sophisticated subjectivism" that can be realist in Enoch's strong, robust

The particularities of Enoch's theory aside, this resolves a tension between belief-dependent theories of obligation and the desire for a realist theory that is also *objective*. Typically, if an object of theorizing is belief-dependent, we are inclined to categorize it as subjective, but the criterion of normative relevance states that if the belief-dependence itself is just a specification of another objective principle (like doxasticism), then the belief-dependence is of no problem for an objective and realist theory.

## 6: OBJECTION FROM VERIDICALITY

The form of dependence doxasticism claims to hold between an agent's moral obligations and their non-moral beliefs is most easily explained by the supervenience relation generally taken to hold between the moral and natural facts.

### *Doxasticist Supervenience*

For every property of having a moral obligation  $M$ , if an agent is  $M$ , then that agent has the property of having a set of relevant non-moral beliefs  $N$  such that necessarily, every agent that is  $N$  is  $M$ .

One challenge that arises from this principle is what I will call the Veridicality Challenge. The Veridicality Challenge stems from the fact that under the principle of doxastic supervenience, the hard facts of the world are entirely trimmed out of the supervenience base. As such, none of the agent's beliefs need be veridical, but there seem to be cases where obligations that stem from non-veridical beliefs will result in morally undesirable outcomes. To illustrate, consider the case of hallucinatory harm.

### *Hallucinatory Harm*

Clare and Dale are in two different worlds and share many natural properties. In fact, they share every natural property concerning

---

sense. For example, Spelman's theory could be one such theory, if the truth of subjectivism itself is not mind-dependent.

the content of their relevant non-moral beliefs. However, Clare and Dale differ in one crucial respect. While most of Clare's non-moral beliefs are veridical, Dale is massively hallucinating; while he believes that he is in the exact same physical situation in his world that Clare believes she is in her world, Dale's differs greatly. Both believe that they see a drowning child, and go to save them, making identical physical motions. Clare successfully saves the child by pulling them forcefully by the arm up and out of the pond, while Dale, making the same motions, harms an innocent bystander relaxing on the ground.

According to the Veridicality Challenge, it seems obvious that Dale and Clare do not share the same moral obligations, and because Dale and Clare share every property concerning the content of their non-moral beliefs, doxastic supervenience must be wrong. I will present and respond to two ways that this denial can be justified. The first I will call the Denial from Moral Competence, and the second I will call the Denial from Impossibility.

According to the Denial from Moral Incompetence, Clare and Dale cannot share the supervenient obligations because Dale is not a morally competent agent. For a person to have a moral obligation, there must be a deeper constitutive sense in which they are obligation-apt, and Dale does not meet whatever constitutive criterion this is. This is further supported by the fact that Dale is not morally responsible and blameworthy for the harm that he causes.

We can begin with an investigation of the relationship between moral responsibility and being obligation-apt. Gary Watson distinguishes between two "faces" of moral responsibility: it has one face of attributability, and one face of accountability (Watson 1996). On the attributability face of responsibility, we merely attribute actions to moral agents insofar as their action is the product of their moral agency. On the accountability face, we adopt a valenced attitude and hold agents responsible as blameworthy or praiseworthy.<sup>5</sup> It is true that Dale is not morally responsible for the harm he causes in the sense that we would not hold him accountable or blameworthy. Dale is also not morally responsible for the harm in the attributive sense; we would not attribute Dale's action to his moral qualities. While Dale is delusional, his delusions are merely sensory delusions; Dale's moral

---

5. See McKenna (2012, 16-17) for a description of neutrally valenced attributive responsibility.

faculties are as sharp as anyone else's. In this way, Dale differs from paradigm cases of moral incompetence, such as Susan Wolf's case of JoJo, a boy whose moral education is corrupted by his dictator parent (Wolf 1987, 53-54). JoJo is a paradigm case of the exculpation of moral responsibility from moral incompetence precisely because his moral faculties are compromised. No such parallel can be drawn from Dale. While Dale, in both senses of moral responsibility, is not responsible for the harm he causes, Dale still is morally responsible for *something*. Dale still tried to save an apparently drowning child, and he is responsible for that attempt; in fact, he may even be praiseworthy to some small extent. We can further specify which sense of responsibility his attempt arises from: Dale's attempt to save a drowning child is the product of his moral agency, and it is this constitutive dimension of moral responsibility that is required for obligation-aptness, so the Denial from Moral Incompetence fails.

According to the Denial from Impossibility, Clare and Dale cannot share the supervenient obligation because Dale cannot possess it, for the reason that there is no drowning child for Dale to be obligated towards. This is no problem for doxasticism; while Clare and Dale do not have obligations identical in referent, they do have obligations identical in content, which is all that is required to have indistinguishable supervenient moral properties. Clare and Dale's obligations are not to the drowning child directly, but rather to the content of their non-moral beliefs that identify an object of obligation. In Clare's case of veridical belief, the object in the obligation can then be further identified with the drowning child, and in Dale's case, the obligation can then be further identified with his hallucination of a child, but there is no reason that these further identifications must be contained within their initial moral properties of having a particular moral obligation.

The Denial from Impossibility encounters a different problem as it commits us to skepticism about our moral obligations. If whether we have a moral obligation depends on there being a world outside of ourselves such that our obligations are to the beings in that world, then our knowledge of our obligations would require us to know the world to which we are obligated. However, this requires defeating traditional skeptical hypotheses such as the possibility that we are brains in vats, dreaming, being deceived by an evil demon, or in this case, merely severely hallucinating. Accepting doxastic supervenience thereby comes with an advantage when compared to views that face this problem; if moral obligations

## compos mentis

supervene merely on non-moral beliefs, then even if we are, say, brains in vats, we can still have moral obligations, and we can still come to know what they are.

## CONCLUSION

What have I argued in this paper? The most important point is that the existing debate over what moral obligations depend on is seriously flawed, because it is missing the view that I take to be right: doxasticism, which occupies the sweet spot of maintaining belief-dependence while not encountering the substantial objections to subjectivism, a quality that each other view in the field lacks. In some sense, doxasticism is simply a lengthy development of the idea that you can only ever act on the information that you have and cannot be obligated to act on the information that you do not.

## REFERENCES

- Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Fogal, Daniel, and Olle Risberg. 2020. "The Metaphysics of Moral Explanations." In *Oxford Studies in Metaethics: Volume 15*, edited by Russ Shafer-Landau, 170-194. Oxford: Oxford University Press.
- Hanson, Louise. 2018. "Moral Realism, Aesthetic Realism, and the Asymmetry Claim." *Ethics* 129 (1): 39-69.
- Jackson, Elizabeth. 2020. "The relationship between belief and credence." *Philosophy Compass* 15 (6): 1-13.
- Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101, no. 3: 461-482.
- Kant, Immanuel. 2007. *Critique of Pure Reason*. London: Palgrave Macmillan.
- King, Alex. 2019. *What We Ought and What We Can*. New York: Routledge.
- Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Wiley-Blackwell.

- McKenna, Michael. 2012. *Conversation and Responsibility*. New York: Oxford University Press.
- Motiz, Mizrahi. 2012. "Does 'Ought' Imply 'Can' from an Epistemic Point of View?" *Philosophia* 40 (4): 829-840.
- Searle, John. 2013. "What is a Speech Act?" In *Philosophy in America*, edited by Max Black, 221-239. New York: Routledge.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1 (3): 229-243.
- Smith, Holly. 2010. "The 'Prospective View' of Obligation." *Journal of Ethics and Social Philosophy* 1: 1-9.
- Spelman, Jonathan. 2017. "Moral Obligation, Evidence, and Belief." PhD diss., University of Colorado Boulder.
- Watson, Gary. 1996. "Two Faces of Responsibility." *Philosophical Topics* 24 (2): 227-248.
- Wedgewood, Ralph. 2013. "Rational 'Ought' Implies 'Can.'" *Philosophical Issues* 23: 70-92.
- Wolf, Susan. 1987. "Sanity and the Metaphysics of Responsibility" In *Responsibility, Character, and the Emotions: New Essays in Moral Psychology*, edited by Ferdinand Schoeman, 46-62. Cambridge: Cambridge University Press.
- Zimmerman, Michael. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge: Cambridge University Press.
- Zimmerman, Michael. 2014. *Ignorance and Moral Obligation*. Oxford: Oxford University Press.